



Cloud technology drives scientific research

Cloud research tools provide a platform for new avenues of scientific research by providing fast access to bare metal resources, writes Robert Roe

Cloud computing offers a more flexible alternative than traditional HPC installations, particularly for scientists and researchers who have varied workloads or that require computing resources to scale with their workloads. Cloud computing can be used to supplement existing cloud resources or enable research for organisations that do not have their own HPC or research computing systems.

In a recent webcast on Scientific Computing World cloud provider Oracle presented a discussion featuring prominent researchers using their platform alongside an Oracle representative. Taylor Newell, senior manager, HPC product management at Oracle who highlighted the Oracle cloud infrastructure and the benefits of leveraging bare metal compute resources to researchers via the cloud.

The webcast focused on the research which drives scientific, technological and social innovation. Focusing on the notion that all too often research can be constrained by the limitations of only using traditionally available computing technologies. Existing technologies may not be the most advanced, they may have capacity thresholds and typically will have a queuing system in place.

These conditions impede progress and impact research outcomes. Cloud technologies, by contrast, open up new possibilities, enabling researchers to access powerful resources and alternative technologies, exactly as and when they are required.

The webcast featured presentations from professor Imre Berger, director of the Max Planck Centre for Minimal Biology at the University of Bristol; Dr Dan Ruderman, director of analytics and machine learning at the Lawrence J. Ellison Institute for Transformative Medicine of the University of South Carolina (USC), USA, and Eric Granthier, head of database services group, CERN IT department at CERN, Switzerland.

Each presenter highlighted their own research objectives from trying to understand the best cancer treatment for a particular patient to creating a synthetic self-assembling ADCOM platform for efficient vaccination or delivering data analytics for massive computing installations at CERN. Each requires massive computing power to deliver scientific progress for their research specialties.

Without access to the cloud these researchers and the organisations they represent would not be able to carry out this research without installing their own cluster or getting access to other organisations' HPC resources.

The cloud allows them to pursue this research quickly and effectively without relying on in-house resources.

During his presentation Newell shared his long history of using and building cloud systems for research purposes. This started with the use of local cloud resources at university and then continued as he moved into the commercial arena.

One thing that Newell stressed was that during this time he and his colleagues were always looking for more compute, more capacity and more storage. Very quickly, that led to launching these types of simulations on the cloud.

"A few years ago I had the opportunity to transition from the consumption side to the product side where we were building the cloud capacity for this type of research," added Newell.

Why are researchers using the cloud? Oracle has historically been a database company which Newell notes "which always needed the best hardware to sit underneath the database software." In 2016 the company extended that service to the public cloud. "Today there are researchers running many different workloads on the public cloud. We wanted to give other people access to this hardware," said Newell.

When this service launched in 2016 there was coverage of just a single region and a few core services across compute, storage, database and networking. Since then, Oracle Cloud has expanded to more



"All of this is being done on bare metal compute behind the scenes"

than 60 services available in 21 cloud regions worldwide with a plan to reach 36 total regions by the end of 2020. In 2019 alone, Oracle Cloud Infrastructure launched more than 200 new services, features, and enhancements.

Oracle Cloud is a Generation 2 enterprise cloud that delivers powerful compute and networking performance and includes a comprehensive portfolio of infrastructure and platform cloud services.

capable of millions of read and write transactions per second.

Oracle's cloud networking services are not over-subscribed, so each tenant gets predictable high-performance and low latency. Based on third-party testing, in a White paper published by Oracle entitled: 'Oracle Cloud Infrastructure Platform Overview' it reported that Oracle and its affiliates offer 2-5 the I/O performance of comparable on-premises or AWS products, with consistent low latency, low jitter and higher bandwidth.

Oracle Cloud Infrastructure is designed for applications that require consistent performance, including raw processing through CPUs or GPUs, millions of storage IOPS, high throughput, and low latency. Performance translates into faster results and greater productivity for end-users.

The platform also uses a "zero-trust architecture". This means that not only are tenants isolated from one another, but tenants are also isolated from Oracle and vice versa. Above Oracle Cloud's core infrastructure are layer upon layer of defenses including encryption everywhere, least-privilege identity and access management, and granular resource and network control all the way out to the edge. Oracle Cloud also uses a strict code security development and deployment processes, a full compliance team that is constantly auditing new regions and services, and a round-the-clock Security Operations Center to guard against threats.

Oracle's Cloud Infrastructure is a natural fit for many high-performance and I/O intensive computing workloads due to this focus on HPC hardware, enabling workloads such as product simulations, risk modeling, and digital twin design and development. These workloads involve huge data sets that need to be analysed using large-scale compute jobs, which demand high performance, high throughput, and low variability. Typical multi-tenant clouds have hypervisor overhead and performance variability. With Oracle's single tenant model, there is no overhead.

Using the cloud also removes the

"My responsibility since I joined Oracle has been to build out the hardware for research and these types of workloads"

long provisioning cycle for acquiring and setting up an HPC cluster in an on-premises environment. Using the cloud researchers can spin up powerful HPC instances in minutes. Moreover, bare-metal instances come with 25 Gbps network throughput, which helps move massive amounts of data quickly.

Oracle also offers HPC-specific instances with higher clock speeds and RDMA-based cluster networking. This provides even higher bandwidth (100 Gbps) and even lower latency (1.5 µs) for HPC workloads that rely on MPI message passing interfaces.

"All of this is being done on bare metal compute behind the scenes" adds Newell. "We scale up and have much different computing requirements than running a web server. This bare metal approach is very important to the simulations and workloads we run in the research space" noted Newell.

"My responsibility since I joined Oracle has been to build out the hardware for research and these types of workloads," added Newell.

Oracle researchers with HPC requirements can access the latest Intel and AMD processors, GPUs from Nvidia based on Pascal and Volta architectures and a high-speed RDMA backed which provides the networking resources required by HPC workloads.

"At the end of the day my responsibility is to make sure that researchers around the world have access to the highest performing hardware in the cloud," Newell concluded.

Each of the individual research projects covered in the Oracle webcast will be highlighted in an upcoming series of articles on Scientific Computing World. Over the next month you can see how these researchers are using cloud to deliver scientific progress and drive their respective fields forward."

If you would like more information regarding this topic or to view the researchers presentations in full the webcast is now available on the Scientific Computing World to be viewed at any time.

www.scientific-computing.com/webcasts